

This Page Is Inserted by IFW Operations  
and is not a part of the Official Record

## **BEST AVAILABLE IMAGES**

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

**IMAGES ARE BEST AVAILABLE COPY.**

**As rescanning documents *will not* correct images,  
please do not report the images to the  
Image Problem Mailbox.**

[> home](#) [> about](#) [> feedback](#) [> login](#)

US Patent &amp; Trademark Office

Try the *new* Portal design

Give us your opinion after using it.

## Search Results

Search Results for: **[throughput and partition and manager and load and balancing and hypervisor ]**

Found 1 of 138,663 searched.

## Search within Results

[> Advanced Search](#)[> Search Help/Tips](#)

---

**Sort by:** Title Publication Publication Date ScoreBinder

---

**Results 1 - 1 of 1** short listing

---

**1** Principles of database buffer management

77%

Wolfgang Effelsberg , Theo Haerder

**ACM Transactions on Database Systems (TODS)** December 1984

Volume 9 Issue 4

This paper discusses the implementation of a database buffer manager as a component of a DBMS. The interface between calling components of higher system layers and the buffer manager is described; the principal differences between virtual memory paging and database buffer management are outlined; the notion of referencing versus addressing of database pages is introduced; and the concept of fixing pages in the buffer to prevent uncontrolled replacement is explained. Three basic t ...

---

**Results 1 - 1 of 1** short listing

---

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.

[> home](#) [> about](#) [> feedback](#) [> login](#)

US Patent &amp; Trademark Office

Try the *new* Portal design

Give us your opinion after using it.

Citation

**ACM Transactions on Database Systems (TODS)** [>archive](#)Volume 9 , Issue 4 (December 1984) [>toc](#)**Principles of database buffer management****Authors**

Wolfgang Effelsberg IBM Scientific Center, Heidelberg, W. Germany

Theo Haerder Univ. Kaiserslautern, Kaiserslautern, W. Germany

**Publisher**

ACM Press New York, NY, USA

Pages: 560 - 595 Periodical-Issue-Article

Year of Publication: 1984

ISSN:0362-5915

<http://doi.acm.org/10.1145/1994.2022> (Use this link to Bookmark this page)[> full text](#) [> abstract](#) [> references](#) [> citings](#) [> index terms](#) [> review](#) [> peer to peer](#)[> Discuss](#) [> Similar](#) [> Review this Article](#)

Save to Binder

[> BibTex Format](#)[↑ FULL TEXT:](#) Access Rules

pdf 2.39 MB

[↑ ABSTRACT](#)

This paper discusses the implementation of a database buffer manager as a component of a DBMS. The interface between calling components of higher system layers and the buffer manager is described; the principal differences between virtual memory paging and database buffer management are outlined; the notion of referencing versus addressing of database pages is introduced; and the concept of fixing pages in the buffer to prevent uncontrolled replacement is explained. Three basic tasks have to be performed by the buffer manager: buffer search, allocation of frames to concurrent transactions, and page replacement. For each of these tasks, implementation alternatives are discussed and illustrated by examples from a performance evaluation project of a CODASYL DBMS.


[↑ REFERENCES](#)



US Patent &amp; Trademark Office

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide



THE ACM DIGITAL LIBRARY

Feedback

Terms used

**workload** and **resource** and **directives** and **operating** and **system** and **partition** and **inter partition** and **manag**
Sort results by Display results 

Save results to a Binder

☐ Search Tips

☐ Open results in a new window
Tr  
Tr

Results 1 - 20 of 200

Best 200 shown

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)**1** [Distributed operating systems](#)

Andrew S. Tanenbaum, Robbert Van Renesse

December 1985 **ACM Computing Surveys (CSUR)**, Volume 17 Issue 4

Full text available: pdf(5.49 MB)

Additional Information: [full citation](#), [abstract](#), [references](#),

Distributed operating systems have many aspects in common with centralized ones, but they also introduction to distributed operating systems, and especially to current university research about distributed operating system and how it is distinguished from a computer network, various key de current research projects are examined in some detail ...

**2** [The state of the art in locally distributed Web-server systems](#)

Valeria Cardellini, Emiliano Casalicchio, Michele Colajanni, Philip S. Yu

June 2002 **ACM Computing Surveys (CSUR)**, Volume 34 Issue 2

Full text available: pdf(1.41 MB)

Additional Information: [full citation](#), [abstract](#), [references](#),

The overall increase in traffic on the World Wide Web is augmenting user-perceived response time with special events. System platforms that do not replicate information content cannot provide th and to match rapid and dramatic changes in the number of clients. The need to improve the perfc of novel content delivery architectures. This article w ...

**Keywords:** Client/server, World Wide Web, cluster-based architectures, dispatching algorithms, mechanisms

**3** [Fast detection of communication patterns in distributed executions](#)

Thomas Kunz, Michiel F. H. Seuren

November 1997 **Proceedings of the 1997 conference of the Centre for Advanced Studies on**

Full text available: pdf(4.21 MB)

Additional Information: [full citation](#), [abstract](#), [references](#),

Understanding distributed applications is a tedious and difficult task. Visualizations based on proc understanding of the execution of the application. The visualization tool we use is Poet, an event i However, these diagrams are often very complex and do not provide the user with the desired ov display repeated occurrences of non-trivial commun ...

**4** [Query evaluation techniques for large databases](#)

Goetz Graefe

June 1993

**ACM Computing Surveys (CSUR)**, Volume 25 Issue 2

Full text available:  pdf(9.37 MB)

Additional Information: [full citation](#), [abstract](#), [references](#),

Database management systems will continue to manage large data volumes. Thus, efficient algorithm sequences will be required to provide acceptable performance. The advent of object-oriented and problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large database systems manipulate simple records, query-processing ...

**Keywords:** complex query evaluation plans, dynamic query evaluation plans, extensible database systems, operator model of parallelization, parallel algorithms, relational database systems, set-n

5 A model for recentralization of computing: (distributed processing comes home)

Harold Lorin

March 1990

**ACM SIGARCH Computer Architecture News**, Volume 18 Issue 1

Full text available:  pdf(1.38 MB)

Additional Information: [full citation](#), [abstract](#), [index terms](#)


Distributed systems commonly contain heterogeneity at all levels of systems structure, differentiation and architecture within a single system. On the other hand, large mainframes tend to be more homogeneous multiprocessors. This paper explores a way of using the models of heterogeneous distributed computing appropriate restructuring of the mainframe can achieve a convolution ...

6 Manageability, availability, and performance in porcupine: a highly scalable, cluster-based mail

Yasushi Saito, Brian N. Bershad, Henry M. Levy

August 2000

**ACM Transactions on Computer Systems (TOCS)**, Volume 18 Issue 3

Full text available:  pdf(2.52 MB)

Additional Information: [full citation](#), [abstract](#), [references](#),

This paper describes the motivation, design and performance of Porcupine, a scalable mail server and scalable electronic mail service using a large cluster of commodity PCs. We designed Porcupine balancing, automatic configuration, and graceful degradation in the presence of failures. Key to the performance is that sessions, data, and underlying ...

**Keywords:** cluster, distributed systems, email, group membership protocol, load balancing, replication

7 Disco: running commodity operating systems on scalable multiprocessors

Edouard Bugnion, Scott Devine, Mendel Rosenblum

October 1997

**ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM**  
Volume 31 Issue 5

Full text available:  pdf(2.30 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [in](#)

8 Manageability, availability and performance in Porcupine: a highly scalable, cluster-based mail

Yasushi Saito, Brian N. Bershad, Henry M. Levy

December 1999

**ACM SIGOPS Operating Systems Review , Proceedings of the seventeenth ACM**  
Volume 33 Issue 5

Full text available:  pdf(1.62 MB)

Additional Information: [full citation](#), [abstract](#), [references](#),

This paper describes the motivation, design, and performance of Porcupine, a scalable mail server and scalable electronic mail service using a large cluster of commodity PCs. We designed Porcupine balancing, automatic configuration, and graceful degradation in the presence of failures. Key to the performance is that sessions, data, and underlying service ...

9 Disco: running commodity operating systems on scalable multiprocessors

Edouard Bugnion, Scott Devine, Kinshuk Govil, Mendel Rosenblum

November 1997

**ACM Transactions on Computer Systems (TOCS)**, Volume 15 Issue 4

Full text available:  pdf(400.76 KB)

Additional Information: [full citation](#), [abstract](#), [references](#),


In this article we examine the problem of extending modern operating systems to run efficiently on a large implementation effort. Our approach brings back an idea popular in the 1970s: virtual machines. We demonstrate our approach with a prototype ...

**Keywords:** scalable multiprocessors, virtual machines

10 Approximate Methods for Analyzing Queueing Network Models of Computing Systems

K. Mani Chandy, Charles H. Sauer

September 1978 **ACM Computing Surveys (CSUR)**, Volume 10 Issue 3


Full text available:  pdf(3.10 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index](#)

11 Cluster-based scalable network services

Armando Fox, Steven D. Gribble, Yatin Chawathe, Eric A. Brewer, Paul Gauthier


October 1997 **ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM**  
Volume 31 Issue 5

Full text available:  pdf(2.42 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index](#)

12 Interposed request routing for scalable network storage

February 2002 **ACM Transactions on Computer Systems (TOCS)**, Volume 20 Issue 1

Full text available:  pdf(363.12 KB)

Additional Information: [full citation](#), [abstract](#), [references](#),


This paper explores interposed request routing in Slice, a new storage system architecture for high block storage. Slice interposes a request switching filter---called a *μproxy*---along each client's network adapter or switch). The *μproxy* intercepts request traffic and distributes it across a server ensemble to handle file service traffic, and explore the ...

**Keywords:** Content switch, file server, network file system, network storage, request redirection

13 Metascheduling for continuous media

David P. Anderson

August 1993 **ACM Transactions on Computer Systems (TOCS)**, Volume 11 Issue 3

Full text available:  pdf(1.64 MB)

Additional Information: [full citation](#), [abstract](#), [references](#),

Next-generation distributed systems will support continuous media (digital audio and video) in that they use continuous media need guaranteed end-to-end performance (bounds on throughput and system components such as CPU schedulers, networks, and file systems must offer performance guarantees). Components, negotiating end-to-end guarantees ...

**Keywords:** multimedia, resource management

14 An architecture for secure wide-area service discovery

Todd D. Hodes, Steven E. Czerwinski, Ben Y. Zhao, Anthony D. Joseph, Randy H. Katz

March 2002 **Wireless Networks**, Volume 8 Issue 2/3

Full text available:  pdf(365.68 KB)

Additional Information: [full citation](#), [abstract](#), [references](#),

The widespread deployment of inexpensive communications technology, computational resources end devices poses an interesting problem for end users: how to locate a particular network service and devices. This paper presents the architecture and implementation of a secure wide-area providers use the SDS to advertise descriptions of available ...

**Keywords:** location services, name lookup, network protocols, service discovery

- 15 Scheduling and resource allocation: SHARP: an architecture for secure resource peering  
Yun Fu, Jeffrey Chase, Brent Chun, Stephen Schwab, Amin Vahdat  
October 2003 **Proceedings of the nineteenth ACM symposium on Operating systems principles**


Full text available:  pdf(339.51 KB)

Additional Information: [full citation](#), [abstract](#), [references](#),

This paper presents Sharp, a framework for secure distributed resource management in an Internet. Sharp is a construct to represent cryptographically protected resource *claims*---promise intervals---together with secure mechanisms to subdivide and delegate claims across a network of flexible *resource peers* ...


**Keywords:** peer-to-peer, resource allocation, resource peering

- 16 Experience Using Multiprocessor Systems—A Status Report  
Anita K. Jones, Peter Schwarz  
June 1980 **ACM Computing Surveys (CSUR)**, Volume 12 Issue 2

Full text available:  pdf(4.48 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

- 17 Serverless network file systems  
Thomas E. Anderson, Michael D. Dahlin, Jeanna M. Neefe, David A. Patterson, Drew S. Roselli, Randolph February 1996 **ACM Transactions on Computer Systems (TOCS)**, Volume 14 Issue 1


Full text available:  pdf(2.69 MB)

Additional Information: [full citation](#), [abstract](#), [references](#),

We propose a new paradigm for network file system design: serverless network file systems. When a server machine, a serverless system utilizes workstations cooperating as peers to provide all file cache, or control any block of data. Our approach uses this location independence, in combination performance and scalability then ...


**Keywords:** RAID, log cleaning, log structured, log-based striping, logging, redundant data storage

- 18 Serverless network file systems  
T. E. Anderson, M. D. Dahlin, J. M. Neefe, D. A. Patterson, D. S. Roselli, R. Y. Wang  
December 1995 **ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM**  
Volume 29 Issue 5

Full text available:  pdf(2.48 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

- 19 A performance analysis of alternative multi-attribute declustering strategies  
Shahram Ghandeharizadeh, David J. DeWitt, Waheed Qureshi  
June 1992 **ACM SIGMOD Record , Proceedings of the 1992 ACM SIGMOD international**  
2

Full text available:  pdf(1.07 MB)

Additional Information: [full citation](#), [abstract](#), [references](#),

During the past decade, parallel database systems have gained increased popularity due to their

characteristics. With the predicted future database sizes and the complexity of queries, the scalability of processors is essential for satisfying the projected demand. Several studies have repeatedly demonstrated that a parallel database system is contingent on the physical characteristics of the hardware.

**20** A client-aware dispatching algorithm for web clusters providing multiple services

Emiliano Casalicchio, Michele Colajanni

April 2001 **Proceedings of the tenth international conference on World Wide Web**




Full text available:  pdf(311.46 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

**Keywords:** clusters, dispatching algorithms, load balancing

Results 1 - 20 of 200

Result page: **1** [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM. [Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)